

From Small Seeds Grow Fruitful Trees: How the PHelpS Peer Help System Stimulated a Diverse and Innovative Research Agenda over 15 Years

J. Vassileva¹ · G. I. McCalla¹ · J. E. Greer¹

Published online: 4 November 2015

© International Artificial Intelligence in Education Society 2015

Abstract PHelpS was a system that helped Correctional Service Canada (CSC) workers to find appropriate helpers among their peers when they were encountering problems while interacting with the CSC database. This seemingly simple system had substantial, and surprising, ramifications. Over time it transformed each of our perspectives as to the issues facing AIED. In this paper we reflect on the influence of the PHelpS peer help system on our subsequent research agenda as well as some of the broader influences of our work. In particular, we discuss a number of research projects arising out of PHelpS directly or indirectly, including the I-Help (aka iHelp) system, a peer help system that has been widely deployed in university courses; a distributed multi-agent architecture for peer help systems that uses fragmented learner modelling to support its activities; the active user modelling paradigm which views “learner model” as a computation not a knowledge structure; the ecological approach, a general architecture for learning systems in which patterns mined from learner interactions with learning objects inform pedagogical decisions; investigations, especially into privacy and reputation, arising from the large scale deployment of iHelp supported by evidence mined from iHelp data; and research into novel affective and social motivation techniques. We conclude by discussing the implications of the common perspective that has emerged from these interrelated research projects. This perspective views the goal of learning technology design to be to track learners as they carry out authentic activities, to deeply understand these learners and their learning context, and to provide just in time support for their learning.

Keywords Peer help · Active learner modeling · Open learner modeling · Ecological approach · Privacy · Reputation · Motivation · Affect · Educational data mining · Agent based learning architectures · Simulation · Scalability

✉ G. I. McCalla
mccalla@cs.usask.ca

¹ ARIES Laboratory, Department of Computer Science, University of Saskatchewan, Saskatoon, SK, Canada

Introduction

From small seeds can grow big trees. When we developed the PHelpS system (Greer et al. 1998a) it seemed like we were tackling a fairly straightforward goal: helping workers at Correctional Service Canada (CSC) as they learned to use a new and complex database system called, rather mystically, OMS – “Offender Management System”. Our big insight in designing PHelpS was to *not* build a training system. Instead, we built a peer help system, where the workers who knew how to use the OMS would help other workers who did not. As we explored how to create such a peer help system, a huge number of interesting issues arose, leading to innovative ideas embedded in PHelpS itself, and broad and deep implications that have effectively informed many of our research activities ever since. This paper first reviews PHelpS and its immediate implications for AIED (the “seed”), and then discusses the large number of interesting projects and ideas that emerged over the next 15 years that are direct outgrowths of the PHelpS system (the “tree”). While each of us has gone in his or her own directions, the ideas we have developed have, in fact, cross-fertilized, and, to stretch the metaphor way too far, have borne interesting fruit.

The Seed: Supporting Synchronous Task-Oriented Peer Help Using PHelpS

PHelpS supported the training needs of Correctional Service Canada (CSC) by helping workers as they carried out real tasks. The key feature of PHelpS was its ability to assist in locating an appropriate peer to help a worker who was having problems while using the OMS to accomplish a task. At the heart of PHelpS were two “minimalist” AI methodologies. One was a knowledge representation scheme capturing at many levels of detail the authentic tasks carried out at CSC. The other was user models, capturing various aspects of each worker, particularly which tasks they had carried out and how well (essentially an overlay on the task hierarchies). If a worker needed help at a particular step in a task, he or she could send a help request to PHelpS, which was able to consult the user models of other workers to recommend potential helpers, i.e. peers who were *ready* (were online at the time of the help request), *willing* (were logged into PHelpS and had indicated availability to help), and *able* (who knew how to accomplish that step and had proved able to help others in the past). The matchmaking functionality of PHelpS was one of the early examples of expertise finding (Yimam-Said and Kobsa 2003), an area that developed further into collaborative enterprise systems (Yu et al. 2011) and more general people recommender systems (Pizzato et al. 2010).

PHelpS would not make the final decision as to who to recommend – the help seeker could choose among the recommended peers, and potential helpers were allowed to choose not to help. This was done by making the user models inspectable (Bull and Pain 1995), an early use of what we in AIED now call “open learner models”. The help seeker could scan through the user models of the potential peer helpers in order to choose somebody who had the specific combination of qualities that they preferred, e.g. a particular pattern of task knowledge, or characteristics such as having the same gender, being of approximately the same age, or being in a similar position (same job title or same union). In addition, any peer helpers who had been contacted could

look at the knowledge profile of the worker needing help and had the right to refuse to help. Finally, workers could inspect their own models to ensure accuracy in the modelling. Open (and scrutable) learner modelling has become a significant trend in the area of AIED since the late 1990s (Dimitrova et al. 1999; Kay 2000; Hartley and Mitrovic 2002). There is a two part special issue of IJAIED on open learner modelling (Dimitrova et al. 2007) and an interesting discussion in (Bull and Kay 2010) of a wide range of uses for open learner modelling and possible future directions for open modelling. Work continues on open learner modelling, especially as it supports learner reflection and other metacognitive skills.

In a proof-of-concept experiment we tested PHelpS with real workers in the Saskatoon Regional Psychiatric Centre of CSC, using a limited number of simulated (but realistic) tasks. PHelpS proved easy to use and mostly recommended good helpers (i.e. helpers who aided the workers in overcoming their problems). Several important issues emerged, particularly the need for the helper to be able to quickly establish a shared context with the help seeker and the usefulness of the task hierarchies in helping the workers to track their own performance as well as to provide the initial basis for this shared context if they needed help. This context problem led to a Ph.D. research project on “helping the peer helper” (Kumar et al. 1999).

As we gained more experience with PHelpS, we began to realize the potential of this system to help make companies into truly “learning organizations” (a buzz word at the time that largely seemed to mean allowing workers to occasionally attend courses). If all workers (including even management) could be committed to helping each other on an ongoing basis, then knowledge would spread throughout the organization at a very deep and situated level. Peer help, we realized, had the potential to literally transform organizations (Ogata and Yano 1999). Peer help supported by technology is now, of course, a big part of many organizations. Further, people are now used to consulting a wide variety of web sources when they need help, and they often also provide information and help on line. So, not only organizations, but also all of us in our everyday lives, are being deeply influenced by the idea of peer help.

However, PHelpS raised several important issues related to knowledge-based peer help systems:

- *knowledge engineering issues*: Task hierarchies need to be verified for their accuracy, need to be updated from time to time as tasks change and to ensure consistency across the different units within a organisation, and may need to be generalized to non-hierarchical task structures (reflecting different organizational cultures). Moreover, there are a host of interesting issues as to what to represent in user models, how to capture user characteristics, how to efficiently manage potentially thousands of user models, what kinds of reasoning are needed/possible over such distributed user models, and so on. Even though the OMS was a fairly straightforward information system, the expense of task and user modelling for an entire organization was daunting.
- *the motivation issue*: A major necessity for the success of systems like PHelpS is to motivate the workers to support each other. The workers know that there are implicit rewards associated with peer help, but in addition the organisation needs to come up with explicit rewards to encourage peer support (e.g. inventing forms of organizational recognition).

- *privacy issues*: The individual user models employed by PHelpS to help workers could also be used by management to monitor workers, perhaps to the workers' perceived or actual disadvantage. Even in organizations where workers are used to being monitored, there may still be major concerns for workers about revealing weaknesses through seeking help online or through providing inadequate help. Worker empowerment and wholehearted worker acceptance are critical to the success of systems like PHelpS.
- *application issues*: While PHelpS was oriented to workplaces with structured tasks, it soon became evident that peer help (and more generally distributed learning systems) could be valuable well beyond such workplaces. This included other types of less structured workplaces and obvious extensions into “traditional” training and educational environments. But this distributed learning approach also seemed to suggest interesting possibilities even in situations requiring “just in time” learning outside any formal educational environment, what we would call today “lifelong learning”.

As it turned out, all of these issues were to be the subject of future research singly or collectively among the three of us. They turned out to be major issues for AIED, too. The rest of this paper will examine some of the research that grew out of the seed that was PHelpS. The PHelpS project itself ended when the sponsoring organization decided that the cost of implementing PHelpS, especially the intensive task modelling needed, was too high. Possibly, as well, they were concerned about the privacy and security issues (important in a correctional service) inherent in using a system that depended so heavily on open user models. Their already planned traditional training approach (which involved courses, travel, human trainers and tutorials), though very expensive as well, seemed to them to be less risky.

From PHelpS to iHelp

Our first “post-PHelpS” foray was to move from the workplace to university learning. By the late 1990's there were already numerous positive examples of implementing on-line course materials and discussion groups at other universities, and commercial course-management software was nascent – CourseInfo LLC/Blackboard (from 1997), Desire2Learn (from 1999), and there were even older enterprise knowledge management tools, such as those by OpenText (dating back to 1991). However, merely providing access to appropriate material via a network does not solve the problem of providing help. One way to decrease the load on teachers is to create conditions in which students help each other. Peer help has many pedagogical advantages. For example, it has the potential to promote the socializing of students in the context of their studies and increase their motivation by giving them social recognition for their knowledge and helpfulness.

In the late 1990's our computer science courses at University of Saskatchewan were overflowing with students, while at the same time our resources to help these students were dwindling. The solution: peer help. Students could help each other for no monetary reward, but the students would gain by attaining deeper insights about the course both through helping and being helped. So, we built a new system, called I-Help

(later iHelp), to be used in our large computer science courses (Greer et al. 1998b). Large classes contained up to 400 students with traditional lectures and laboratory activities. Each course was considered to be a closed community of students, tutors, markers, and professor(s). One subsystem of I-Help, the “1-on-1” component, drew directly from PHelpS to provide a tool that would find helpers for people who encountered problems in a course. These helpers were chosen from among other people in the same course community. The other component, building on techniques from another early peer help system called CPR (Cooperative Peer Response) (Bishop et al. 1997), was an open peer forum where students in a course community could post comments or questions and receive responses from each other. In later versions of iHelp there was also a chat room to support synchronous interaction within each course community.

iHelp was used by thousands of students over a 10 year period ending when some of the now standard course management tools started to fulfill many of the iHelp functions. Over these years, many experiments were run (Bull et al. 2001) generating a massive amount of data. One important finding was that instructors had a good deal of influence on student engagement in iHelp. Instructors who strongly promoted the use of the tool, who were visibly present in the online discussions, and who encouraged the peer help model, seemed to engender greater engagement by their students. Part of this effect was to increase student motivation and part was to inspire student confidence that iHelp was a legitimate form of academic support. Another finding was that student engagement was not uniform. Many students were quiet lurkers, who read postings but contributed nothing. A small minority of students seemed to enjoy the “status” associated with being a visible helper and the notoriety of being an “expert” among their peers. Both of these findings confirmed that motivation was a key issue in peer help systems, and led to further research in our laboratory (as discussed later in this paper).

iHelp made a big difference to many students. Some claimed that the iHelp discussion forum was the most important factor in their success in particular university courses. The peer helper recommender (the 1-on-1 component of iHelp - a direct outgrowth of PHelpS) wasn't used as much as we had hoped due to the cost of building task models, the lack of a critical mass of users required for synchronous help sessions, and the need to deal better with motivational issues. Both the successes and failures of iHelp led to many interesting research questions, to innovative methodologies, and to original architectures for supporting learning. Some of these will be discussed next.

A Distributed Multi-Agent Architecture for Peer Help Systems

In the design of the I-Help system for courses, we wanted to map the distributed nature of the collaboration that was taking place during a peer-help session onto a decentralized software architecture. We were strongly attracted (Vassileva 1998) by the multi-agent system paradigm that was gaining popularity at the time with its metaphor of software components as independent autonomous agents (Nwana 1996) pursuing their own goals, using their own resources, and forming relationships among each other. This led to the MAGALE (Multi-AGENT Adaptive Learning Environment) architecture underpinning the I-Help system (Vassileva et al. 1999).

In MAGALE each student had a personal agent, a novel kind of pedagogical agent that wasn't a "third party" agent (such as a learning companion (Chan 1996)), but which actually represented the student's own interests in the learning system. To do this, the personal agent kept a model of the user's competences, preferences and relationships (like a "friends list" in Facebook terminology, although online social networks did not exist yet at that time). The personal agent provided a semi-anthropomorphic interface, an avatar (an animated cartoon figure) chosen by the student to represent him or her when requesting and receiving help. The personal agent represented the student's interests (both as helpers and help seekers) in the negotiations when helpers were being selected for a help request. A facet of these negotiations involved determining a payment amount for help, in I-Help Credit Units, a virtual currency that students earned as a "salary" for using I-Help as well as in payment from other students for helping them out. The currency was introduced as a potential solution to the motivation issue identified in PHelpS. In this way an "I-Help economy" was created that was supposed to regulate the supply and demand of help (as a traded valuable commodity) during course "crunch" times such as assignment deadlines. In addition to personal agents, MAGALE had special agents representing various online resources as well as matchmaker agents that used various matching criteria and algorithms to find suitable peers.

The number of user models in MAGALE was actually much higher than the number of agents. This was because each agent was responsible for gathering information about any other personal agent with which that agent interacted, and since each personal agent was a "proxy" of a given user, indirectly the agent model reflected the model of that user. Thus each user had many models about him- or herself, stored by many different personal agents with whom his/her agent had negotiated and the user had interacted with (Vassileva 2001). This was a natural encapsulation of responsibilities in the distributed MAGALE architecture and resulted in higher overall flexibility in the system compared to an aggregated centralized repository of user models. The resulting "fragmentation" of the user models throughout the system had both positive and negative sides: negative, since a centralized solution is usually more efficient computationally; positive, since keeping the user model with the agent that created it allowed preserving the context (the help request, constraints, negotiation preferences) and thus, potentially, providing richer information for recommending helpers and building groups of similar users.

This multi-agent architecture was easily extendible and allowed for natural distribution of resources. Social interaction among agents also provided a rich metaphor that suggested further functionalities and a realistic context for investigating theoretical issues in the area of multi-agent systems including models of bilateral negotiation (Winoto et al. 2002, 2005), interpersonal relationships (trust) (Breban and Vassileva 2002; Wang and Vassileva 2003), and formation of communities of agents with similar preferences (Wang and Vassileva 2004) acting as "neighbourhoods" in recommender systems. These issues were becoming very "hot" in the areas of multi-agent and peer-to-peer systems as well as distributed sensor networks. Our work influenced a lot of other research in these areas, resulting in many hundreds of citations and influencing the work of Endriss (2006); Zhou and Hwang (2007); Boukerch et al. (2007); Moyaux et al. (2006); Liu et al. (2006) and many others.

The MAGALE architecture was the basis for the I-Help 1-on-1 component, which was deployed over 4 consecutive terms between the Fall of 1999 and the Spring of 2001 (Greer et al. 2001). In these deployments there were over 1000 personal agents (representing over 1000 different learners) and many application and matchmaking agents. Each user ended up with up to 20 “fragmented” models of himself/herself, including the model held by their own personal agent, of course, but also models held by several matchmaker agents, as well as other users' personal agents which had interacted with the user's personal agent. In total there were over 10,000 such fragmented models distributed through the I-Help system. The information kept in each model contained preferences, rankings, ratings, and numeric overlays on course topics depending on which agent created the model and for what purpose. The key to making sense of the distributed user models was the ability to interpret multi-modal information from multiple heterogeneous relevant sources and to integrate this information as needed into a user model of appropriate granularity. The main questions boiled down to how to manage all this information:

- how to locate the agent that has a model of relevant user characteristics, given the context and the purpose for which the model is needed?
- how to make sense of inconsistent and even contradictory user information?
- how to interpret models created by other agents?

In fact, it became evident that when user models are fragmented and distributed in such a decentralized architecture, there is no single user model at all, but merely a subset of user characteristics that are important to the user modelling task at hand. In some real sense the user model ceases to be an object, but becomes a context-based calculation that is carried out when needed for a particular purpose: “model” as a verb, not a noun. This led to our proposal for a new user modelling paradigm called “active user modelling”.

Active User Modelling

The “active user modelling” paradigm was first proposed during the heart of the experiments with the MAGALE architecture (McCalla et al. 2000), and then was elaborated more fully in (Vassileva 2002) and (Vassileva et al. 2003) upon reflection after the MAGALE experiments had come to an end. User modelling in this paradigm is viewed as a process – a computation (function) over a space of four major dimensions: subjects, objects, purposes and resources. For a given user modelling activity, the subject is the person or agent doing the modelling, the object is the person or agent being modelled, the purpose is the adaptation or the activity for which the model is being created, and the resources are the computational or information resources contributing to the modelling process. For a typical I-Help 1-on-1 help request, for example, the subject might be the personal agent of the student requesting help, the object might be the student themselves, the purpose might be to find an appropriate helper, and the resources would be other agents like the personal agents of potential helpers, matchmaker agents, etc. The active user modelling paradigm shifts the focus from traditional AI and user modelling issues of *representation* (consistency,

representation schemes, indexing) to the *process* of collecting, interpreting and utilising user data for a particular purpose. In short, the active modelling paradigm is an example of a procedural, rather than a declarative, approach (Rumelhart and Norman 1983) and makes specific claims as to the kinds of processes and contextualizations necessary for user modelling.

In order to understand how to build such active modelling systems, we implemented a multi-agent recommender system for investment portfolio management (Niu et al. 2004), which delivers increasingly better recommendations depending on the availability of resources (both referees and computational). At its core was a library of purpose-based user modelling processes that could be pre-selected by the designer to be combined and executed at run time depending on the context and resource availability. This was one of the first approaches in the area of user modelling that considered the problem of reusing user modelling processes. It influenced work on user modelling standards (Paramythis and Loidl-Reisinger 2004), ubiquitous user modelling (Heckmann 2006), generic user models (Kobsa 2007), and user model interoperability (Dolog and Schäfer 2005; Carmagnola et al. 2011). This is an area that has gained a lot of attention recently (Viviani et al. 2010).

The Ecological Approach: An Architecture for Fragmented Learning Systems

I-Help's ideas led naturally to a new view of learning as it might take place in a technology-saturated world, a world where people are deeply and pervasively immersed in technology, a world of constant change demanding continuous ongoing adaptation. In such a world, it is not just the user modelling that is fragmented; learning, teaching, culture, and technology are themselves all fragmented (McCalla 2000). The learner sits at the heart of their own personal "electronic village" that filters their perspectives of, and interactions with, the vast information and social space around them. In such a world learning can often happen opportunistically in the context of other activities, can take advantage of the vast amount of information available on the web, and can be supported by other people who are often members of on-line communities whose membership overlaps the learner's own communities. Technology to support such learning must be accessible "just in time", and the learning must be contextualized by learner goals, the content being learned, personal characteristics, and social factors. The "active modelling" paradigm discussed above certainly synchronizes with this perspective in that it focuses on the same contextual elements, but it is not enough: there is still a need to devise an appropriate architecture for the rest of the learning system. The result was the "ecological approach" (EA) architecture (McCalla 2004).

In the ecological approach, learning environments are abstracted as learning object repositories (Brooks et al. 2006; Richards et al. 2002). Learning object repositories in the EA are very generally defined as including learning material (such as web pages, e-books, talk slides, videos, etc.), learning technology (such as simulation environments, open forums, intelligent tutoring systems, etc.), and other learners. Later versions of the EA, in fact, "activated" these objects as agents (similar to the MAGALE agents) that are responsible for the interactions between the object represented by the agent and

other agents or the outside world. Such agents included personal agents representing learners, tutors, etc. (as in I-Help), as well as agents representing each of the other kinds of learning objects. The EA also assumes that when an agent is interacting with another agent, all interactions are recorded at a fine-grained level (even keystrokes) and stored with the object. In particular the interactions between a learner and a learning object are stored with the learning object, constituting fragmented learner models (as in I-Help) distributed throughout the repository.

In this architecture, then, a learner can explore (or be guided to) appropriate learning objects when trying to fulfill a learning goal. Since the repository keeps a record of their interactions with each object, the learner leaves an e-trail (Driver and Clarke 2008). Over time, with many learners fulfilling many learning goals (or trying to at least), a huge amount of data can be collected about each learner and about each learning object. This data can then be “mined” for various patterns that would fulfill various pedagogical goals. One such goal, explored by (Tang and McCalla 2005), is to recommend an appropriate next learning object to a learner. Basically, such a recommender system looks at the path taken by the learner, abstracts the learner’s patterns of interaction with the learning objects on that path, compares these patterns with the patterns of other learners who have followed a similar path, extracts learners with similar patterns, and then recommends learning objects that have proven successful for these similar learners. Success can be determined easily if the learning object has some way of testing the learner directly, but it may also be possible to infer success (or failure) by observing particular patterns in learner interactions with a learning object. Such educational recommender systems have become an important area of advanced learning technology research (Manouselis et al. 2011).

Other possible goals are to explore, through comparing patterns of interaction stored with various learning objects, which kinds of objects have worked well for which kinds of learners in which situations; to recommend other learners whose behaviour indicates they would be helpful to a learner facing an impasse (as in the I-Help 1-on-1 component); to infer learner models of a more traditional sort (including possibly cognitive, content, and social attributes) extracted from patterns of interaction. In other words, the EA is a general architecture allowing a wide variety of educational environments to be modelled and a large number of pedagogical goals to be achieved. All of these goals anticipate a central role for educational data mining, a field that itself was only just emerging at the time the EA architecture was proposed, but is now a very active area of learning technology research.

Increasingly, the EA has been seen as a good fit for lifelong learning, a burgeoning subfield of AIED (Kay 2008), where learners can be supported in important learning goals as they arise in their lives, and where fine-grained data is continuously collected about learners over the long haul and mined for useful patterns: truly a “big data” problem. An important question has been how to test lifelong learning systems. One answer, at least for initial pre-deployment testing, is simulation. In a seminal paper (VanLehn et al. 1994) first proposed simulation as important for AIED, but the bulk of subsequent work in the field has been on simulated pedagogical agents. Simulation for testing AIED systems and exploring pedagogical questions is just now becoming important, with both high fidelity models such as SimStudent (Matsuda et al. 2007) and low fidelity models as in (Champaign et al. 2011). In our lab we have embarked on a lifelong learning project (Lelei and McCalla 2015) to build a simulated graduate

school environment (that will run for 10 simulated years) in which to explore how various kinds of peer help can affect graduate students.

The EA architecture is now also being tested through simulation experiments, largely low fidelity ones. The first implementation of the EA architecture, the “EA platform” (EAP), explored various strategies for recommending learning objects to learners to determine which strategy worked best (Erickson et al. 2013). Two other simulation experiments have also been carried out in the EA architecture. One of these looked at the effect on learners of observing the performance of their peers (an open modelling issue) (Frost and McCalla 2013). Another experiment explored data-driven instructional planning within the EA with a goal of coming up with instructional strategies that would work in dynamic open-ended learning object repositories where new learning objects are constantly being incorporated and old learning objects deleted (Frost and McCalla 2015). While simulation will never replace human testing, it is becoming clear that simulation will increasingly be part of the advanced learning technology designer’s toolkit going forward, especially for low fidelity testing and parameter “tweaking” in the preliminary design phases (as in our ecological approach work).

The ecological approach also has echoes in the emergence of learning services architectures. Our own work on active learning modelling led to an architecture called MUMS (Brooks et al. 2004) that had producers (monitoring learner actions, say) providing information for consumers (such as a learning system component) that processed this information in the way most appropriate to them and their current circumstances. More recently, the Caliper Framework (IMS Global Learning Consortium 2015) provides for “activity metrics” to be computed from interaction data captured as learners take part in specific learning activities. Such learning services frameworks are likely to play an increasingly important role in how we build and maintain advanced learning systems in a technologically heterogeneous and data saturated world.

The Issues of Large Scale Deployment

In addition to the 1-on-1 component, the other big I-Help sub-system, the discussion component, also led to interesting follow-up research. The emphasis was on scalability. One of the advantages of peer help in large university classes is to overcome the problems of scale. In courses with hundreds or even thousands of students (such as MOOCs), a peer help solution becomes not only scalable, but affordable. A new version of the I-Help system, re-engineered after the MAGALE experiments and relabelled iHelp, was deployed in most of the computer science courses at the University of Saskatchewan, and other courses as well. The discussion component in iHelp was refined into a large scale open peer forum suitable for thousands of users with a flexible model for managing users and groups with sophisticated identity management, role-based access control, and detailed user/usage tracking. This discussion forum became the basis for early research into visualization in social networks (Brooks et al. 2007), identity management and pseudonymous identity tracking (Richardson and Greer 2004), and privacy-enhanced reputation management (Anwar and Greer 2012). These are important topics for advanced learning technology. Social network analysis, including visualization, has become increasingly useful in online

learning (Carolan 2013), and there are numerous challenges in cyber-security that must be overcome to achieve safe, scalable, online learning environments Bandara et al. (2014).

As a platform for research the discussion forums of iHelp also provided a means to mine discussion forum data for useful patterns, to find frequently asked questions and answers, and to examine the semantic annotation of discussion threads. The rich fine-grained data about learner activities, queries, conversations, and helping behaviours led us to examine more closely several semantic web issues (Brooks et al. 2009). The discussion forum data remains a valuable resource that continues to be data mined for interesting patterns.

Motivating Participation

Even with the many deployments of I-Help, one of the hardest problems we encountered, particularly in the 1-on-1 component, was a low rate of student participation (Vassileva et al. 2001), (Vassileva 2012). We used virtual currency as an incentive for students to answer questions, yet the question of how to “cash” the virtual currency in for something meaningful for the students was not really solved – students cared for grades, but we could not (ethically) give course grades in exchange for participation in an experimental system. Therefore, we started looking into the use of affective and social motivators, rather than tangible extrinsic rewards. One approach (Okonkwo and Vassileva 2001) investigated the persuasive impact of an animated pedagogical agent displaying emotions to learners in response to their learning performance. The results showed that, while not contributing any significant performance gain in learning, the incorporation of emotional feedback changed the way students perceived the learning process, making it more engaging. We also found gender-based and individual differences in the user perception of an emotional agent, which need to be taken into account when designing more adaptive and “intelligent” emotional pedagogical agents. This work influenced further research into the effect of emotional feedback on learners by Corradini et al. (2005); Kim and Baylor (2006), and Yee et al. (2007). In fact affective feedback in learning has become a very active research area (Beale and Creed 2009), (Vail et al. 2015). Nearly a decade later we continued this stream of work now in the context of persuasive technologies design (Hamari et al. 2014). We moved from design-based research to large scale studies of the impact of personality (in this case, gamer type, but also gender) on the effect of persuasive strategies on promoting healthy eating (Orji et al. 2014).

Social Learning Environments

Another direction of research followed up on Slashdot’s successful motivational strategy by allowing users to build reputation in peer-to-peer online learning communities. This strategy makes sense only in social learning environments (Vassileva 2008), where learners form communities across time and space to

learn together. A peer-to-peer system called Comtella (Vassileva 2002; Vassileva 2004) was designed and deployed for sharing course-related resources contributed by users. User participation was measured and rewarded with points, which were awarded at different user levels, associated with different powers/rights, and interfaces (Bretzke and Vassileva 2003; Cheng and Vassileva 2006). Visualizing the levels and contributions of all the users in a group turned out to be a very effective motivator for the active users, who engaged in social comparison (Sun and Vassileva 2006; Vassileva and Sun 2007). The social visualization idea was developed further by us and others (Farzan and Brusilovsky 2011; Farzan et al. 2008) to motivate users in various social applications, e.g. social networks, recommender systems and discussion forums (Webster & Vassileva, 2006; Sahib & Vassileva, 2009; Raghavun & Vassileva, 2011). The idea of using social visualization and social comparison to promote reflection and participation was developed further by Sharon Hsiao and Peter Brusilovsky into Social Open Learner Modeling (Hsiao, Bakalov, Brusilovsky & König-Ries, 2011).

We developed a dynamic adaptive reward mechanism that calculated the points awarded for different user actions depending on how useful for the community this type of action was at the current moment and depending on historical data about the quality of participation of the individual user. The mechanism was incorporated into Comtella and evaluated during one term in a class on Ethics and IT, and the results showed that it regulated in a sustainable way the contributions in an online community (Cheng & Vassileva, 2006). This means that there was active participation exactly to a desired level predefined by the instructor for timing, quantity and quality of contributions. Participants did not over-contribute nor were there substandard (fake) contributions. This is because the incentive mechanism effectively encouraged contributions that were timely (early in the week) and not too abundant (a personal limit was set for each student depending on his/her previous quality of contribution). To determine in advance the appropriate quantitative rewards for any new deployment we used a system dynamics model to simulate an online community and predict the changes in participation based on the timing and size of rewards (Mao, Vassileva & Grassmann, 2007). Our research on incorporating incentives for participation in the system design had significant impact (hundreds of Citations of the publications about Comtella) and anticipated the full emergence of the area of gamification (Deterding et al., 2011; Kapp, 2012).

Implications and Lessons Learned

These various research projects coalesce around a perspective on learning that involves fine-grained tracking of learners as they carry out authentic activities, deep understanding of these learners (on many dimensions), and just in time support for them in various ways as they learn. This perspective is highly tuned to a world with a vast amount of information and a large number of learners who have access to very good technology and who have a great need for learning across many aspects of their lives. In other words, our world!

With the recent explosion of interest in MOOCs, learning analytics, etc., the world has come to understand what AIED researchers have known all along: that online learning will be a critically important technology going forward, a technology that has the potential to dramatically transform teaching and learning. MOOCs still have a scalability problem, however. While the content of the course (video lectures, online material, web pages, quizzes, etc.) can be consumed by an essentially unlimited number of learners at a time, each learner has idiosyncratic needs while learning. Tending to these needs does not scale well, since it requires personalization of a high calibre that cannot be met in a massive course by a small cadre of human instructors. Fortunately, such personalization has been one of the main concerns of AIED over the 40+ years of the field's existence, so there is much research to draw on. In particular, we feel that the PHelpS/I-Help experience provides two big insights: (i) use the learners themselves as helpers; and (ii) support these learners through personalization technology that can find the right helper at the right time in clever and scalable ways. Our work, along with that of many other researchers in AIED and related advanced learning technology fields, should find a receptive application context going forward.

The fruitful tree that has grown out of the PHelpS seed also illustrates that it is important for researchers to continue to explore ideas that may at first seem fairly uninteresting. Starting with the seemingly prosaic goal of helping workers use a new database system, our investigations have led directly and indirectly to a huge number of interesting research projects that have influenced (and, of course, been influenced by) AIED, AI, data mining, multi-agent systems, social computing, semantic web, etc. We will not repeat specific contributions here (since they have been mentioned in the appropriate sections of this paper), but among the AIED areas that our research has most impacted have been learner modelling (especially open learner modelling), personal agents, learning communities, motivation, privacy, simulation, distributed e-learning architectures, educational data mining, and lifelong learning. We think there is still much more to explore in all of these areas, and more implications of our research to come.

Another important lesson is to look at how the system being designed will really be used. For example, at the very beginning of the PHelpS design process we had thought that workers would need help with actually using the OMS database, but we observed that they frequently asked questions about issues involving a bigger organizational context, such as subtleties about the kinds of information to enter into various database fields, rather than how to use various features of the company software. It is also important to be very open minded going forward and to pursue issues as they arise – new research ideas have a habit of organically growing out of questions that have come up in previous research, often in unpredictable directions. Finally, while consulting the research literature is crucial, it is also critical for researchers to directly interact with other researchers as projects evolve. Synergies that arise in conversations among researchers can lead to much deeper and more interesting ideas. Such interactions can be internal (with colleagues and graduate students) or external (with researchers from other labs). Research itself is a learning experience with a large social component.

Acknowledgments We are grateful for the funding we have received over the years from the Natural Sciences and Engineering Research Council of Canada through Discovery Grants to each of the authors and through other grants. We would also like to thank our many graduate students, undergraduate students, research assistants, and postdocs who have been a constant source of good ideas and hard work.

References

- Anwar, M., & Greer, J. E. (2012). Facilitating trust in privacy-preserving E-learning environments. *IEEE Transactions on Learning Technologies.*, 5(1), 62–73.
- Bandara I., Ioras F. & Maher I K. (2014). Cyber Security Concerns in E-Learning Education. In *Proceedings of ICERI2014 Conference, 17th-19th November, 7 p.*
- Beale, R., & Creed, C. (2009). Affective interaction: how emotional agents affect users. *International Journal of Human-Computer Studies*, 67(9), 755–776.
- Bishop A. S., Greer J.E. & Cooke J.E. (1997). The Co-operative Peer Response System: CPR for Students. *Proceedings of ED-MEDIA 1997, AACE, June, 172–178.*
- Boukerch, A., Xu, L., & EL-Khatib, K. (2007). Trust-based security for wireless ad hoc and sensor networks. *Computer Communications*, 30, 2413–2427.
- Breban S. & Vassileva, J. (2002). Using Inter-Agent Trust Relationships for Efficient Coalition Formation, R. Cohen & B. Spencer (eds.) *Proceedings of the 13th Canadian Conference on AI*, Calgary, 28–30 May, 2002, Springer Verlag LNAI 2338, 221–236.
- Bretzke, H., & Vassileva, J. (2003). Motivating cooperation in peer to peer networks, *User Modeling UM03*, Johnstown, PA, June. Springer Verlag LNCS, 2702, 218–227.
- Brooks C., Winter M., Greer J.E. & McCalla G.I. (2004). The massive user modelling system (MUMS), *7th International Conference on Intelligent Tutoring Systems (ITS-04)*, Maceió, Brazil, August 2004, 635–645.
- Brooks, C., Bateman, S., McCalla, G. I., & Greer, J. E. (2006). Applying the agent metaphor to learning content management systems and learning object repositories, 8th International Conference on Intelligent Tutoring Systems (ITS-06). Taiwan, June: Jhongli, 808–810.
- Brooks, C., Greer, J. E., & Parchoma, G. (2007). Understanding learning communities, *Workshop on Assessment of Group and Individual Learning Through Intelligent Visualization*, at AIED 2007, California, 6 p.
- Brooks C., Bateman S., Greer J.E. & McCalla G.I. (2009). Lessons learned using social and semantic web technologies for E-learning, Chapter in *Semantic Web technologies for e-learning*, Dicheva, Mizoguchi, Greer eds. Amsterdam, IOS Press, 262–280.
- Bull, S., & Kay, J. (2010). Open learner models. Berlin Heidelberg: In *Advances in Intelligent Tutoring Systems*. Springer, 301–322.
- Bull, S., & Pain, H. (1995). *Did I say what I think I said, and do you agree with me?: inspecting and questioning the student model*. DAI Research Paper: University of Edinburgh.
- Bull S., Greer J.E., McCalla G.I. & Kettel L. (2001). User modelling in I-Help: what, why, when and how, *8th International Conference on User Modeling (UM 2001)*, Sonthofen, Germany, 117–126.
- Carmagnola, F., Cena, F., & Gena, C. (2011). User model interoperability: a survey. *User Modeling and User-Adapted Interaction*, 21(3), 285–331.
- Carolan, B. V. (2013). *Social Network Analysis and Education: Theory, Methods & Applications*. Sage Publications, Thousand Oaks, California, 344 p.
- Champaign J., Zhang J. & Cohen R. (2011). Coping with poor advice from peers in peer-based intelligent tutoring: the case of avoiding bad annotations of learning objects. *User Modeling, Adaption and Personalization*, Springer Berlin Heidelberg, 38–49.
- Chan, T. W. (1996). Learning companion systems, social learning systems, and the global social learning club. *Journal of Artificial Intelligence in Education*, 7(2), 125–159.
- Cheng, R., & Vassileva, J. (2006). Design and evaluation of an adaptive incentive mechanism for sustained educational online communities. *User Modelling and User-Adapted Interaction*, 16(2/3), 321–348.
- Corradini, A., Mehta, M., BERNSEN, N., & Charfuelan, M. (2005). Animating an interactive conversational character for an educational game system. *Proceedings Intelligent User Interfaces IUI, 05*, 183–190.
- Deterding S., Sicart M., Nacke L., O'Hara K. & Dixon D. (2011). *Gamification. Using Game-Design Elements in Non-Gaming Contexts*. *Proceedings of CHI 11*, ACM, 2425–2428.

- Dimitrova, V., Self, J., & Brna, P. (1999). The interactive maintenance of open learner models, proc. In *9th International Conference on Artificial Intelligence in Education*, Le Mans, France 405–412.
- Dimitrova V., McCalla G.I. & Bull S. (2007). Guest editors, special issue on open learner modelling of the *International Journal of Artificial Intelligence in Education*, 17, 2 and 17, 3, April and July 2007.
- Dolog P & Schäfer M (2005). A Framework for Browsing, Manipulating and Maintaining Interoperable Learner Profiles. *Proc. User Modeling 2005*, Springer, 397–401.
- Driver, C., & Clarke, S. (2008). An application framework for mobile, context-aware trails. *Pervasive and Mobile Computing*, 4, 719–736.
- Endriss U. (2006). Monotonic Concession Protocols for Multilateral Negotiation. *Proceedings of the 5th International Joint Conference on Autonomous Agents and Multiagent Systems AAMAS '06*, 392–299.
- Erickson G., Frost S., Bateman S. & McCalla G.I. (2013). Using the Ecological Approach to Create Simulations of Learning Environments. *Proceedings of 16th Int. Conference on Artificial Intelligence in Education (AIED 2013)*, Memphis, July, 411–420.
- Farzan, R., & Brusilovsky, P. (2011). Encouraging user participation in a course recommender system: an impact on user behavior. *Computers in Human Behavior*, 27(1), 276–284.
- Farzan R., DiMicco J.M., Millen D. R., Dugan C., Geyer W. & Brownholtz E.A. (2008). Results from deploying a participation incentive mechanism within the enterprise. *Proceedings of the SIGCHI Conference on Human Factors in Computing System 2008s*, 563–572.
- Frost, S., & McCalla, G. I. (2013). Exploring through simulation the effects of peer impact on learning, AIED workshop on simulated learners, part 4. In *Proceedings of the workshops at the 16th International Conference on Artificial Intelligence in Education (AIEDWS 2013)*, Vol 1009, Memphis, TN, July, 21–30.
- Frost S. & McCalla G. I. (2015). Exploring Through Simulation an Instructional Planner for Dynamic Open Ended Learning Environments. *Proc. 17th International Conference on Artificial Intelligence in Education (AIED 2015)*, Madrid, Spain, 4 p.
- Greer, J. E., McCalla, G. I., Collins, J., Kumar, V., Meagher, P., & Vassileva, J. (1998a). Supporting peer help and collaboration in distributed workplace environments. *International Journal of Artificial Intelligence in Education*, 9, 159–177.
- Greer J.E., McCalla G.I., Cooke J.E., Collins J., Kumar V., Bishop A. & Vassileva J. (1998b). The Intelligent helpdesk: Supporting Peer-Help in a University Course. *Proc. Intelligent Tutoring Systems Conference (ITS 1998)*, Springer Berlin, 494–503.
- Greer J.E., McCalla G.I., Vassileva J., Deters R., Bull S. & Kettel L. (2001). Lessons Learned in Deploying a Multi-Agent Learning Support System: The I-Help Experience. *Proc. of AI in Education Conference AIED 2001*, San Antonio, IOS Press: Amsterdam, 410–421.
- Hamari, J., Koivisto, J., & Pakkanen, T. (2014). Do persuasive technologies persuade? - a review of empirical studies. In *Persuasive Technology: Lecture Notes in Computer Science Volume, 8462(2014)*, 118–136.
- Hartley, D., & Mitrovic, A. (2002). Supporting learning by opening the student model. In *Proc. 6th international conference on Intelligent Tutoring Systems 2002, Lecture Notes in Computer Science Volume 2363* (pp. 453–462).
- Heckmann, D. (2006). *Ubiquitous user modeling*. IOS Press, Amsterdam, 268 p.
- Hsiao, I. H., Bakalov, F., Brusilovsky, P., & König-Ries, B. (2011). Open social student modeling: visualizing student models with parallel introspective views. In *Proc. 19th International Conference on User Modeling, Adaptation and Personalization*, Springer 171–182.
- IMS Global Learning Consortium (2015). Learning Measurement for Analytics Whitepaper, <http://imsglobal.org/IMSLearningAnalyticsWP.pdf>, 12 p.
- Kapp, K. M. (2012). *The Gamification of Education*. John Wiley & Sons, Hoboken, New Jersey, 336 p.
- Kay J. (2000). Stereotypes, Student Models and Scrutability. *Intelligent Tutoring Systems*, Springer Berlin Heidelberg, 19–30.
- Kay, J. (2008). Lifelong learner modeling for lifelong personalized pervasive learning. *IEEE Transactions on Learning Technologies*, 1(4), 215–228.
- Kim, Y., & Baylor, A. (2006). A social-cognitive framework for pedagogical agents as learning companions. *Educational Technology Research & Development*, 54(6), 569–596.
- Kobsa A. (2007). Generic User Modeling Systems. *The Adaptive Web*. LNCS 4321, Springer, 136–154.
- Kumar, V., McCalla, G. I., & Greer, J. E. (1999). Helping the peer helper. In *Proc. 9th International Conference on Artificial Intelligence in Education*, Le Mans, France, IOS 325–332.
- Lelei E. & McCalla G.I. (2015). Using Simulation to Explore Reciprocal Help Seeking in a Lifelong Learning Context. *Int. J. of Business Process Integration and Management*, 7 (3), Inderscience, 228–246.
- Liu, K., Bhaduri, K., Das, K., Nguyen, P., & Kargupta, H. (2006). Client-side web mining for community formation in peer-to-peer environments. *Proc. WebKDD 2006. SIGKDD Explorations*, 8(2), 11–20.

- Manouselis N., Drachsler H., Vuorikari R., Hummel H. & Koper R. (2011). Recommender systems in technology enhanced learning. In *Recommender Systems Handbook*. Springer, 387–415
- Mao Y., Vassileva J. & Grassmann W. (2007). A System Dynamics Approach to Study Virtual Communities, in *Proc. IEEE HICSS'07 Mini-Track on Virtual Communities*. Big Island, Hawaii, January, 187a.
- Matsuda N., Cohen W.W., Sewall J., Lacerda G., & Koedinger K.R. (2007). Predicting Students Performance with SimStudent that Learns Cognitive Skills from Observation, *Proc. Int. Conf. on Artificial Intelligence in Education*, Marina del Rey, 467–476.
- McCalla, G. I. (2000). The fragmentation of culture, learning, teaching and technology: implications for the artificial intelligence in education research agenda in 2010. *International Journal of Artificial Intelligence in Education*, 11(2), 177–196..
- McCalla G.I. (2004). The Ecological Approach to the Design of E-Learning Environments: Purpose-based Capture and Use of Information about Learners. *J. of Interactive Media in Education (JIME)*, 2004 (1), 18 p. [<http://www-jime.open.ac.uk/2004/7/mccalla>]
- McCalla G.I., Vassileva J., Greer J.E. & Bull S. (2000). Active Learner Modelling. *Intelligent Tutoring Systems Conference ITS 2000*, Montreal, pp. 53–62
- Moyaux, T., Chaib-Draa, B., & D'Amours, S. (2006). Supply chain management and multiagent systems: an overview. *Studies in Computational Intelligence (SCI)*, 28, 1–27.
- Niu, X., McCalla, G. I., & Vassileva, J. (2004). Purpose-based expert finding in a portfolio management system. *Computational Intelligence Journal*, 20(4), 548–561.
- Nwana, H. (1996). Software agents: an overview. *Knowledge Engineering Review*, 11(3), 1–40.
- Ogata, H., & Yano, Y. (1999). Combining social networks and collaborative learning in distributed organizations. *Proc. EdMedia*, 1999, 119–125.
- Okonkwo, C., & Vassileva, J. (2001). Affective pedagogical agents and user persuasion. In *Proceedings of the 9th international conference on human-computer interaction, new Orleans* (pp. 397–401).
- Orji, R., Vassileva, J., & Mandryk, R. (2014). Modeling the efficacy of persuasive strategies for different gamer types in serious games for health. *User Modeling and User Adapted Interaction (UMUAI)*, 24(5), 453–498.
- Paramythis, A., & Loidl-Reisinger, S. (2004). Adaptive learning environments and e-learning standards. *Electronic Journal on e-Learning*, 2(1), 181–194.
- Pizzato L., Rej T., Chung Th., Koprinska I. & Kay J. (2010). RECON: a reciprocal recommender for online dating, *International Conference on Recommender Systems RecSys 2010*, September, Barcelona, Spain, 207–214.
- Raghavun, K., & Vassileva, J. (2011). Visualizing reciprocity to motivate participation in an online community. In *Proc. 5th IEEE International Conference on Digital Ecosystems and Technologies Conference (DEST)*, IEEE Press (pp. 89–94).
- Richards, G., McGreal, R., Hatala, M., & Friesen, N. (2002). The evolution of learning object repository technologies: portals for on-line objects for learning. *International Journal of E-Learning and Distance Education*, 17(3), 67–79.
- Richardson, B.R., & Greer, J.E. (2004). An architecture for identity management. *Privacy Security and Trust Conference*, Fredericton, August, 103–108.
- Rumelhart, D. E., & Norman, D. A. (1983). *Representation in Memory*, Center for Human Memory Technical Report: University of San Diego, 117 p.
- Sahib, Z., & Vassileva, J. (2009). Designing to attract participation in a niche community for women in science and engineering, *proc. In Workshop on Social Computing in Education, 1st IEEE International Conference on Social Computing SocialComp 2009*. Vancouver: August.
- Sun L. & Vassileva J. (2006). Social Visualization Encouraging Participation in Online Communities. *Groupware: Design, Implementation, and Use*, *Proc. 12th International CRIWG Workshop*, Lecture Notes in Computer Science 4154, Springer-Verlag, Berlin Heidelberg, 349–363.
- Tang, T. Y., & McCalla, G. I. (2005). Paper Annotations with Learner Models, *12th International Conference on Artificial Intelligence in Education AIED-05* (pp. 654–661). Amsterdam: July.
- Vail A., Boyer K., Wiebe E. & Lester J. (2015). The Mars and Venus Effect: The Influence of User Gender on the Effectiveness of Adaptive Task Support. *Proc. User Modeling, Adaptation and Personalization, UMAP'2015*, July, Dublin, 265–276.
- VanLehn, K., Ohlsson, S., & Nason, R. (1994). Applications of simulated students: an exploration. *Int. J. of Artificial Intelligence in Education*, 5, 135–175.
- Vassileva, J. (1998). Goal-based autonomous social agents supporting adaptation and teaching in a distributed environment, *Intelligent Tutoring Systems ITS 1998, San Antonio, LNCS 1452* (pp. 564–573). Berlin: Springer Verlag.

- Vassileva J. (2001). Distributed user modelling for universal information access, *Proc. of the 9th International Conference on Human-Computer Interaction*. New Orleans, Lawrence Erlbaum: Mahwah, N.J., 122–126.
- Vassileva, J. (2002). Supporting peer-to-peer user communities. In R. Meersman, Z. Tari, et al. (Eds.), *On the move to meaningful internet systems 2002: CoopIS, DOA, and ODBASE, Coordinated International Conferences Proceedings*, Irvine, LNCS 2519. Berlin-Heidelberg: Springer Verlag, 230–247.
- Vassileva, J. (2004). Harnessing P2P power in the classroom, *Intelligent Tutoring Systems Conference ITS 2004. Lecture Notes in Computer Science 3220*, 305–314.
- Vassileva, J. (2008). Towards social learning environments. *IEEE Transactions on Learning Technologies*, 1(4), 199–214.
- Vassileva, J. (2012). Motivating participation in social computing applications: a user modeling perspective. *User Modeling and User-Adapted Interaction*, 22(1–2), 177–201.
- Vassileva J. & Sun L. (2007). Using Community Visualization to Stimulate Participation in Online Communities. *e-Service Journal*, 6, 1, 3–40.
- Vassileva J., Greer J.E., McCalla G.I., Deters R., Zapata D., Mudgal C. & Grant S. (1999). A Multi-Agent Approach to the Design of Peer-Help Environments, in *Proc. Int. Conf. on Artificial Intelligence in Education AIED 1999*, Le mans, France, 38-45.
- Vassileva J., Deters R., Greer J.E., McCalla G.I., Bull S. & Kettel L. (2001). Lessons from Deploying I-Help. *Workshop on Multi-Agent Architectures for Distributed Learning Environments at AIED 2001*. San Antonio, 3–11.
- Vassileva, J., McCalla, G. I., & Greer, J. E. (2003). Multi-agent multi-user modeling. *User Modeling and User-Adapted Interaction*, 13(1), 179–210.
- Viviani, M., Bennani, N., & Egedy-Zsigmond, E. (2010). A survey on user modeling in multi-application environments. *Proc. IEEE Advances in Human-Oriented and Personalized Mechanisms, Technologies and Services (CENTRIC)*, 111–116.
- Wang, Y., & Vassileva, J. (2003). Bayesian network-based trust model, *proc. In of IEEE/WIC International Conference on Web intelligence (WI 2003), Halifax, Canada* (pp. 372–378).
- Wang Y. & Vassileva J. (2004). Trust-based community formation in peer-to-peer file sharing networks, *Proc. of IEEE/WIC/ACM International Conference on Web Intelligence (WI 2004)*, Beijing, 341–348.
- Webster A.S. & Vassileva J. (2006). Visualizing Personal Relations in Online Communities, in *Adaptive Hypermedia and Adaptive Web-Based Systems*, Dublin, Ireland, June, Springer LNCS 4018, 223–233.
- Winoto P., McCalla G.I. & Vassileva J. (2002). An Extended Alternating-Offers Bargaining Protocol for Automated Negotiation in Multi-Agent Systems. *Proc. On the Move to Meaningful Internet Systems 2002: CoopIS, DOA, and ODBASE*. 179-194.
- Winoto, P., McCalla, G. I., & Vassileva, J. (2005). Non-monotonic-offers bargaining protocol. *Journal of Autonomous Agents and Multi-Agent Systems*, 11(1), 45–67.
- Yee, N., Bailenson, J., & Rickertsen, K. (2007). A meta-analysis In *Recommender Systems Handbook*. Springer In *Recommender Systems Handbook*. Springer is of the impact of the inclusion and realism of human-like faces on user experiences in interfaces. *Proc. of CHI, 2007*, 1–10.
- Yimam-Said, D., & Kobsa, A. (2003). Expert-finding systems for organizations: problem and domain analysis and the DEMOIR approach. *Journal of Organizational Computing and Electronic Commerce*, 13(1), 1–24.
- Yu, H. T., Liu, C. R., & Zhang, F. Z. (2011). Reciprocal recommendation algorithm for the field of recruitment. *Journal of Information & Computational Science*, 8(16), 4061–4068.
- Zhou, R., & Hwang, K. (2007). PowerTrust: a robust and scalable reputation system for trusted peer-to-peer computing. *IEEE Transactions on Parallel and Distributed Systems*, 18(4), 460–473.